arXiv:1407.7357v1 [cs.IR] 28 Jul 2014

Text Classification Using Association Rules, Dependency Pruning and Hyperonymization

Yannis Haralambous and Philippe Lenca

Institut Mines Telecom, Telecom Bretagne, UMR CNRS 6285 Lab-STICC
Technopôle Brest Iroise CS 83818, 29238 Brest Cedex 3, France
<surname>.<name>@telecom-bretagne.eu

Abstract. We present new methods for pruning and enhancing itemsets for text classification via association rule mining. Pruning methods are based on dependency syntax and enhancing methods are based on replacing words by their hyperonyms of various orders. We discuss the impact of these methods, compared to pruning based on tfidf rank of words.

Introduction

Automatic text classification is an important text mining task, due to the huge number of text documents that we have to manage daily. Text classification has a wide variety of applications such as Web document and email classification. Indeed, most of the Web news services daily provide a large number of articles making them impossible to be organized manually [14]. Automatic subject classification [9] and SPAM filtering [18] are two additional examples of the interest of automatic text classification.

Automatic text classification can be defined as below. Given a set of documents such that each document is labeled with a class value, learn a model that assigns a document with unknown class to one or more particular classes. This can also be done by assigning a probability value to each class or by ranking the classes.

A wide variety of classical machine learning techniques have been used for text classification. Indeed, texts may be represented by word frequencies vectors, and thus most of the quantitative data methods can be used directly on the notorious "bag-of-words" model (cf. [27,3]).

Choosing a classifier is a multicriteria problem. In particular one has often to make a trade-off between accuracy and comprehensibility. In this paper, we are interested in both criteria with a deeper interest in comprehensibility. We are thus interested in rule-based approaches and especially in class association rules algorithms. Several studies have already successfully considered association rule-based approaches in text mining (e.g., [4], [29], [7], [25]). This framework is suitable for considering some statistical characteristics (e.g., high-dimensionality, sparsity...) of the bag-of-words model where a document is represented as a set of words with their associated frequency in the document.

However a text is more than a set of words and their frequencies. Enhancing the bag-of-words approach with linguistic features has also attracted several works (e.g., [12,11,13,22], [23,16,10], [21,6]).

We here propose a class association rules based approach enriched by linguistic knowledge. The paper is organized as follows: after introducing the techniques we are going to use (class association rules § 1.1, dependencies § 1.2, hyperonymization § 1.3) we describe our main algorithms (for training § 2.1, classifying § 2.2 and evaluating § 2.3); follows the experimental section, where we give results obtained by tfidf pruning § 3.2, dependency-based pruning § 3.3 and hyperonymization § 3.4, and, finally, we end up by a conclusion and perspectives for future work § 4.

1 Proposed model for text classification

Let a *corpus* be a set $\mathbf{C} = \{D_1, \ldots, D_n\}$ of documents. Let $\mathcal{C}$ be a set of classes. An *annotated corpus* is a pair $(\mathbf{C}, \text{class})$ where $\text{class} : \mathbf{C} \to \mathcal{C}$ is a function that maps each document D_i to a (predefined) class of $\mathcal{C}$.

A document $D \in \mathbf{C}$ is a set of sentences S. The corpus $\mathbf{C}$ can be considered as a set of sentences $\mathbf{S} = \{S_1, \ldots, S_m\}$ if we go through the forgetful functor (which forgets the document to which the sentence belongs). Repeated sentences in the same document, or identical sentences in different documents are considered as distinct, i.e., there is a function $\iota \colon \mathbf{S} \to \mathbf{C}$ which restores the forgotten information. We extend the class function to $\mathbf{S}$ by $\text{class}(S) := \text{class}(\iota(S))$.

A sentence S is a sequence of words w (sometimes we will consider S simply as a set, without changing the notation). Let $\mathcal{W} = \bigcup_{S \in \mathbf{S}} \bigcup_{w \in S} \{w\}$ be the set of all words of $\mathbf{C}$.

1.1 Class association rules and text classification

Let $\mathcal{I}$ be a set of objects called *items* and $\mathcal{C}$ a set of classes. A *transaction* T is a pair $(\{i_1, \ldots, i_n\}, c)$, where $\{i_1, \ldots, i_n\} \subseteq \mathcal{I}$ and $c \in \mathcal{C}$. We denote by $\mathcal{T}$ the set of transactions, by $\text{items}(T)$ the set of items (or "itemset") of T and by $\text{class}(T)$ the class of T.

Let I be an itemset. The *support* of I is defined by

$$\text{supp}(I) := \frac{\#\{T \in \mathcal{T} \mid I \subseteq \text{items}(T)\}}{\#\mathcal{T}}.$$

Let $\sigma \in [0, 1]$ be a value called *minimum support*. An itemset I is called *frequent* if its *support* exceeds σ.

The *confidence* of a transaction t is defined as

$$\text{conf}(t) := \frac{\#\{T \in \mathcal{T} \mid \text{items}(t) \subseteq \text{items}(T) \wedge \text{class}(t) = \text{class}(T)\}}{\#\{T \in \mathcal{T} \mid \text{items}(t) \subseteq \text{items}(T)\}}.$$

Let $\kappa \in [0, 1]$ be a value called *minimum confidence*. A *class association rule* (or "CAR") $r = (\{i_1, \ldots, i_n\}, c)$ [15] is a transaction with frequent itemset and a *confidence* exceeding κ.

To classify text with CARs, we consider words as being items, documents as being itemsets and pairs of documents and classes as being transactions. The advantage of this technique is that CARs can be easily understood and hence potentially improved by the user, especially if the classifier is tuned so that it produces humanly reasonable number of rules. Once the classifier is trained, to classify a new sentence we first find all CARs whose items are contained in the sentence, and then use an aggregation technique to choose a predominant class among those of the CARs we found.

An important issue of CARs is that the complexity is exponential with respect to the itemset size, and hence we need to keep it bounded in specific ranges, independently of the size of documents to classify. Using entire documents as transactions is computationally out of reach, therefore pruning techniques play an important rôle. Our approach consists in (a) restricting CARs to the sentence level, (b) prune sentences by using morphosyntactic information (cf. § 1.2) and modifying itemsets using semantic information (cf. § 1.3).

1.2 Itemset pruning using dependencies

One can prune sentences either by using word frequencies (cf. § 3.2) or by using information obtained by morphosyntactic parsing (cf. § 3.3). In this paper we introduce the latter approach, in the frame of dependency grammar.

Dependency grammar [28,19] is a syntactic theory, alternative to *phrase-structure analysis* [8] which is traditionally taught in primary and secondary education. In phrase-structure syntax, trees are built by grouping words into "phrases" (with the use of intermediate nodes NP, VP, etc.), so that the root of the tree represents the entire sentence and its leaves are the actual words. In dependency grammar, trees are built using solely words as nodes (without introducing any additional "abstract" nodes). A single word in every sentence becomes the *root* (or *head*) of the tree. An oriented edge between two words is a *dependency* and is tagged by a representation of some (syntactic, morphological, semantic, prosodic, etc.) relation between the words. For example in the sentence "John gives Mary an apple," the word "gives" is the head of the sentence and we have the following four dependencies:

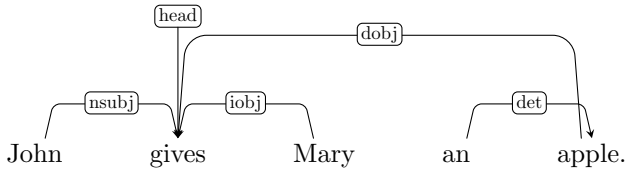

where tags nsubj, dobj, iobj, det denote "noun subject," "direct object," "indirect object" and "determinant."

Let S be a sentence and $\mathcal{D}$ be the set of dependency tags: {nsubj, ccomp, prep, dobj, ... } A *dependency* is a triple (w_1, w_2, d) where $w_1, w_2 \in S$ and $d \in \mathcal{D}$. Let $\mathrm{Dep}(S)$ denote the set of dependencies of S and $\mathrm{root}(S)$ the head of S.

Pruning will consist in defining a *morphosyntactic constraint* ϕ i.e. a condition on dependencies (and POS tags) of words, the fulfillment of which is necessary for the word to be included in the itemset.

But before describing pruning algorithms and strategies, let us first present a second technique used for optimizing itemsets. This time we use semantic information. We propose to replace words by their hyperonyms, expecting that the frequencies of the latter in the itemsets will be higher than those of the former, and hence will improve the classification process.

1.3 Hyperonymization

The WordNet lexical database [20] contains sets of words sharing a common meaning, called *synsets*, as well as semantic relations between synsets, which we will use to fulfill our goal. More specifically, we will use the relations of *hyperonymy* and of *hyperonymic instance*. The graph having synsets as nodes, and hyperonymic relations as edges, is connected and rooted: starting with an arbitrary synset, one can iterate these two relations until attaining a sink. Note that in the case of nouns it will invariably be the synset `00001740` {entity} while for verbs there are approx. 550 different verb sinks.

Let $\mathbb{W}$ be the WordNet lexical database, $s \in \mathbb{W}$ a synset and $h : \mathbb{W} \to 2^{\mathbb{W}}$ the hyperonymic or hyperonymic instance relation. We define an *hyperonymic chain* $\mathrm{CH}(s)$ as a sequence $(s_i)_{i \geq 0}$ where $s_0 = s$ and $s_i \in h(s_{i-1})$, for all $i \geq 1$. Hyperonymic chains are not unique since a given synset can have many hyperonyms. To replace a word by the most pertinent hyperonym, we have to identify the most significant hyperonymic chains of it.

The *wn-similarity* project [24] has released synset frequency calculations based on various corpora. Let $\mathrm{lf}(s)$ denote the logarithmic frequency of synset s in the BNC English language corpus [2] and let us arbitrarily add infinitesimally small values to the frequencies so that they become unique ($s \neq s' \Rightarrow \mathrm{lf}(s) \neq \mathrm{lf}(s')$). We use frequency as the criterion for selecting a single hyperonymic chain to represent a given synset, and hence define the *most significant hyperonymic chain* $\mathrm{MSCH}(s)$ as the hyperonymic chain $(s_i)_{i \geq 0}$ of s such that $s_i = \arg\max_{s \in h(s_{i-1})} \mathrm{lf}(s)$, for all $i \geq 1$. The chain $\mathrm{MSCH}(s)$ is unique thanks to the uniqueness of synset frequencies.

Our CARs are based on words, not synsets. Hence we need to extend MSCHs to words. Let w be a lemmatized word. We denote by $\mathrm{Synsets}(w) \subset \mathbb{W}$ the set of synsets containing w. If the cardinal $\#(\mathrm{Synsets}(w)) > 1$ then we apply a standard disambiguation algorithm to find the most appropriate synset s_w for w in the given context. Then we take $(s_i)_i = \mathrm{MSCH}(s_w)$ and for each synset s_i in this chain we define $h_i(w) = \mathrm{proj}_1(s_i)$ $(i > 0)$, that is the projection of s_i to its first element, which by WordNet convention is the most frequent word in the synset. The function vector $h_* : \mathcal{W} \to \mathcal{W}$ (with $h_0 \equiv \mathrm{Id}$) is called *hyperonymization*, and $h_i(w)$ is the *i-th order hyperonym of w*.

Algorithm 1: Training

Data: An annotated corpus $\mathbf{C}$, values of minimum support σ and minimum confidence κ

Result: A set of CARs $\mathcal{R} = (\{R_1, \ldots, R_N\}, \text{conf})$ where $\text{items}(R_i) \subset \mathcal{W}$, $\text{class}(R_i) \in \mathcal{C}$, and $\text{conf}(R_i)$ is the confidence of rule R_i

`Train(`$\mathbf{C}$, σ, κ`):`

- $\mathbf{S} :=$ forgetful($\mathbf{C}$); $\mathbf{S}' := \emptyset$;
- **for** $S \in \mathbf{S}$ **do**
  - $S' :=$ `Hyperonymize` (`Prune` (`Lemmatize` (S)));
  - $\text{class}(S') :=$ $\text{class}(\iota(S))$;
  - $\mathbf{S}' := \mathbf{S}' \cup \{S'\}$;
- **end**
- $\mathcal{R} :=$ `Apriori` $(\mathbf{S}', \sigma, \kappa)$;

end

2 Operational implementations for document classification

Our text classifier operates by first training the classifier on sentences and then classifying the documents by aggregating sentence classification. These two procedures are described in Sections 2.1 and 2.2 respectively. Specific evaluation procedure is presented in Section 2.3.

2.1 Training

The `Train` algorithm (cf. Alg. 1) takes as input an annotated corpus $\mathbf{C}$ and values of minimum support σ and minimum confidence κ. It returns a set of CARs together with their confidence values.

The first part of the algorithm consists in processing the corpus, to obtain efficient and reasonably sized transactions. Three functions are applied to every sentence:

1. `Lemmatize` is standard lemmatization: let $\mathcal{P}$ be the set of POS tags of the *TreeTagger* system [26] (for example, NP stands for "proper noun, singular", VVD stands for "verb, past tense", etc.), and let $\mathcal{W}'$ be the set of lemmatized forms of $\mathcal{W}$ (for example, "say" is the lemmatized form of "said"); then we define $\lambda : \mathcal{W} \to (\mathcal{W} \cup \mathcal{W}') \times \mathcal{P}$, which sends a word w to the pair (w', p) where w' is the lemmatized form of w (or w itself, if the word is unknown to *TreeTagger*) and p is its POS tag.
2. `Prune` is a function which prunes the lemmatized sentence so that only a small number of (lemmatized) words (and POS tags) remains. Several sentence pruning strategies will be proposed and compared (cf. § 3.2 and 3.3).
3. `Hyperonymize` is a function which takes the words in the pruned itemset and replaces them by the members of their most significant hyperonymic chains. Several strategies will also be proposed and compared (cf. § 3.4).

Algorithm 2: Classification

Data: A set of CARs $\mathcal{R}$, a document D_0
Result: The predicted class predclass(D_0), variety β, dispersion Δ
`Classify(`$\mathcal{R}$, D_0`):`

- **for** $S \in D_0$ **do**
  - **if** $\exists r \in \mathcal{R}$ *such that* items(r) $\subset S$ **then**
    - $R_S := \underset{r \in \mathcal{R} \wedge \text{items}(r) \subset S}{\arg\max} \text{conf}(r)$;
  - **end**
- **end**
- $\text{predclass}(D_0) := \underset{c \in \mathcal{C}}{\arg\max} \sum_{\substack{S \in D_0 \\ \text{class}(R_S)=c}} \text{conf}(R_S)$;
- $\beta := \#\{c \in \mathcal{C} \mid (\text{class}(R_S) = c) \wedge (\text{conf}(R_S) > 0)\}$;
- $\Delta := \max_{c \in \mathcal{C}} \sum_{\substack{S_i \in D_0 \\ \text{class}(R_{S_i})=c}} \text{conf}(R_{S_i}) - \min_{c \in \mathcal{C}} \sum_{\substack{S_i \in D_0 \\ \text{class}(R_{S_i})=c}} \text{conf}(R_{S_i})$;

end

The second part of Alg. 1 uses the *apriori* algorithm [5] with the given values of minimum support and minimum confidence and output restrictions so as to generate only rules with item $c \in \mathcal{C}$ in the consequent. It returns a set $\mathcal{R}$ of CARs and their confidence.

It should be noted that this algorithm operates on individual sentences, hereby ignoring the document level.

2.2 Classification

The `Classify` algorithm (cf. Alg. 2) uses the set of CARs produced by `Train` to predict the class of a new document D_0 and furthermore provides two values measuring the quality of this prediction: variety β and dispersion Δ.

The first part of the algorithm takes each sentence S of the document D_0 and finds the most confident CAR that can be applied to it (i.e., such that the itemset of the rule is entirely contained in the itemset of the sentence). At this stage we have, for every sentence: a rule, its predicted class and its confidence.

Our basic unit of text in `Train` is sentence, therefore CARs generated by Alg. 1 produce a class for each sentence of D_0. An aggregation procedure is thus needed in order to classify the document. This is done by taking class by class the sum of confidence of rules and selecting the class with the highest sum.

Although this simple class-weighted sum decision strategy is reasonable, it is not perfect and may lead to wrong classification. This strategy will be *optimally* sure and robust if (a) the number of classes is minimal, and (b) the values when summing up confidence of rules are sufficiently spread apart. The degree of fulfillment of these two conditions is given by the parameters *variety* β (the number of classes for which we have rules), and *dispersion* Δ (the gap between the most confident class and least confident one). These parameters will contribute to comparison among the different approaches we will investigate.

Algorithm 3: Evaluation

Data: An annotated corpus $\mathbf{C}$, initial values of minimal support σ_0 and confidence κ_0, standard number of rules ρ_0

Result: Values of average precision $\overline{P}$, recall $\overline{R}$, F-measure $\overline{F}$. Values of average number of rules ρ, variety β and dispersion Δ

`SingleEvaluate(`$\mathbf{C}$, σ, κ`):`

- $(\mathbf{C}_1, \ldots, \mathbf{C}_{10})$:= `Partition (Shuffle (`$\mathbf{C}$`),10)`;
- `/* tenfold cross validation */`
- **for** $I \in \{1, \ldots, 10\}$ **do**
  - $(\mathcal{R}_I, \beta_I, \Delta_I)$:= `Train (`$\mathbf{C} \setminus \mathbf{C}_1$, σ, κ`)`;
  - **for** $D \in \mathbf{C}_I$ **do**
    - predclass(D) := `Classify (`$\mathcal{R}_I$, D`)`;
  - **end**
  - **for** $c \in \mathcal{C}$ **do**
    - $R_I(c) := \frac{\#\{d \in \mathbf{C}_I | (\text{predclass}(d)=c) \wedge (\text{class}(d)=c)\}}{\#\{d \in \mathbf{C}_I | \text{class}(d)=c\}}$;
    - $P_I(c) := \frac{\#\{d \in \mathbf{C}_I | (\text{predclass}(d)=c) \wedge (\text{class}(d)=c)\}}{\#\{d \in \mathbf{C}_I | \text{predclass}(d)=c\}}$; $F_I(c) := \frac{2R_I(c)P_I(c)}{R_I(c)+P_I(c)}$;
  - **end**
- **end**
- **for** $c \in \mathcal{C}$ **do**
  - $(R(c), P(c), F(c)) := \frac{1}{10}\sum_{I=1}^{10}(R_I(c), P_I(c), F_I(c))$;
- **end**
- $(\rho, \beta, \Delta) := \frac{1}{10}\sum_{I=1}^{10}(\#\mathcal{R}_I, \beta_I, \Delta_I)$;
- $(\overline{R}, \overline{P}, \overline{F}) := \frac{1}{\#\mathcal{C}}\sum_{c \in \mathcal{C}}(R(c), P(c), F(c))$;

end

`Evaluate(`$\mathbf{C}$, σ_0, κ_0, ρ_0`):`

- (σ, κ) := `FindOptimal`$(\mathbf{C}, \sigma_0, \kappa_0, \rho_0)$;
- $(\overline{R}, \overline{P}, \overline{F}, \rho, \beta, \Delta)$:= `SingleEvaluate`$(\mathbf{C}, \sigma, \kappa)$;

end

2.3 Evaluation

We evaluate the classifier (Alg. 3), by using 10-fold cross validation to obtain average values of recall, precision, F-measure, variety and dispersion. This is done by algorithm `SingleEvaluate`, once we specify values of minimal support and minimal confidence.

Comparing rule-based classification methods is problematic because one can always increase F-measure performance by increasing the number of rules, which results in overfitting them. To avoid this phenomenon and compare methods in a fair way, we fix a number of rules ρ_0 (we have chosen $\rho_0 = 1{,}000$ in order to produce a humanly reasonably readable set of rules) and find values of minimal support and confidence so that F-measure is maximal under this constraint.

Function `FindOptimal` will launch `SingleEvaluate` as many times as necessary on a dynamic grid of values (σ, κ) (starting with initial values (σ_0, κ_0)), so that, at the end, the number of rules produced by `Train` is as close as possible to ρ_0 (we have used $\#\mathcal{R} \in [\rho_0 - 2, \rho_0 + 2]$) and $\overline{F}$ is maximal.

Algorithm 4: Tfidf-based corpus pruning

Data: An annotated corpus (considered as a set of sentences) $\mathbf{S}$
Result: The pruned corpus $\mathbf{S}'$
`Prune(`$\mathbf{S}$, N`):`
  $\mathbf{S}' := \emptyset$;
  for $S \in \mathbf{S}$ **do**
    for $w \in S$ **do**
      $\texttt{Tfidf}_S(w) := \text{freq}_S(w) \cdot \log\left(\frac{\#\{S \in \mathbf{C}\}}{\#\{S \in \mathbf{C} | w \in S\}}\right)$
    end
    $S' := \emptyset$; $S_0 := S$;
    for $i \in \{1, \ldots, N\}$ **do**
      $w' := \arg\max_{S_0} \texttt{Tfidf}_S(w)$;
      $S_0 := S_0 \setminus \{w'\}$; $S' := S' \cup \{w'\}$;
    end
    $\mathbf{S}' := \mathbf{S}' \cup \{S'\}$;
  end
end

3 Experimental results on Reuters corpus

In this section, we investigate three methods: (a) pruning through a purely frequentist method, based on tfidf measure (§ 3.2); (b) pruning using dependencies (§ 3.3); (c) pruning using dependencies followed by hyperonymic extension (§ 3.4).

3.1 Preliminaries

In the Reuters [1] corpus we have chosen the 7 most popular topics (GSPO = sports, E12 = monetary/economic, GPOL = domestic politics, GVIO = war, civil war, GDIP = international relations, GCRIM = crime, law enforcement, GJOB = labor issues) and extracted the 1,000 longest texts of each.

The experimental document set is thus a corpus of 7,000 texts of length between 120 and 3,961 words (mean 398.84, standard variation 169.05). The texts have been analyzed with the Stanford Dependency Parser [17] in collapsed mode with propagation of conjunct dependencies.

3.2 Tfidf-based corpus pruning

Tfidf-based corpus pruning consists in using a *classical* `Prune` function as defined in Alg. 4. It will be our baseline for measuring performance of dependency- and hyperonymy-based methods.

Note that this definition of the tfidf measure diverges from the legacy one by the fact that we consider not documents but sentences as basic text units. This is because we compare tfidf-generated CARs to those using syntactic information, and syntax is limited to the sentence level. Therefore, in order to obtain a

fair comparison, we have limited term frequency to the sentence level and our "document frequency" is in fact a sentence frequency.

Having calculated $\texttt{Tfidf}_S(w)$ for every $w \in S \in \mathbf{S}$, we take N words from each sentence with the highest tfidf values, and use them as transaction items. The performance of this method depends on the value of N. On Fig. 1 the reader can see the values of three quantities as functions of N:

1. F-measure: we see that F-measure increases steadily and reaches a maximum value of 83.99 for $N = 10$. Building transactions of more than 10 words (in decreasing tfidf order) deteriorates performance, in terms of F-measure;
2. variety: the number of predicted classes for sentences of the same document progressively increases but globally remains relatively low, around 3.1, except for $N = 12$ and $N = 13$ where it reaches 4.17;
3. dispersion: it increases steadily, with again a small outlier for $N = 12$, probably due to the higher variety obtained for that value of N.

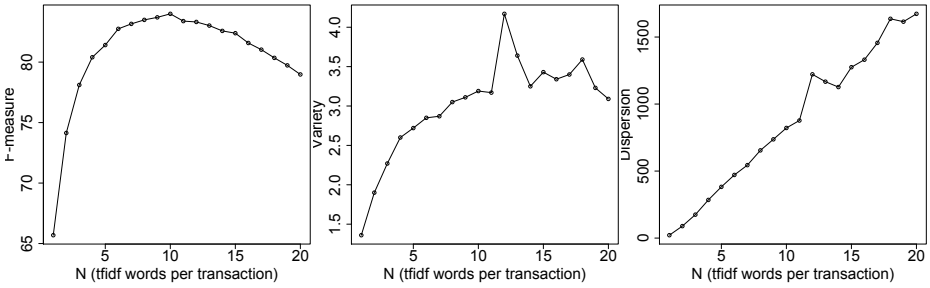

Fig. 1. F-measure, variety and dispersion of tfidf-based pruning methods as a function of the number of words kept in the corpus

Furthermore, each investigated method will generate transactions of various sizes. It is fair to compare them with tfidf-based methods with similar transactions sizes. Therefore we will use the results displayed in Fig. 1 to compare the performance of subsequent methods with the one of the tfidf-based method of similar transaction size. Table 1 presents the results obtained by applying the tdfif-based pruning method, with a single word per transaction ($N = 1$).

Table 1. Tfidf-based pruning, keeping a single word per transaction

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	69.30	44.48	55.44	45.75	52.54	82.90	67.98	59.77
Precision	70.09	77.81	71.25	79.76	71.62	80.78	73.35	74.95
F-measure	69.69	56.60	62.36	58.15	60.61	81.83	70.56	65.69

MinSupp=0.006, MinConf=67.6, Var.=1.36, Disp.=21.53, AvgTransSize=1.00

3.3 Methods based on dependencies

In this section we investigate several strategies using the dependency structure of sentences. Our general approach (cf. Alg. 5) keeps only words of S that fulfill

Algorithm 5: Dependency-based corpus pruning

Data: An annotated corpus **S** and a morphosyntactic contraint
ϕ: $S \to \{\text{true}, \text{false}\}$
Result: The pruned corpus $\mathbf{S}'$
`Prune(`**S**`,` ϕ`):`
  $\mathbf{S}' := \emptyset$;
  for $S \in \mathbf{S}$ **do**
    $S' := \emptyset$;
    for $w \in S$ **do**
      if $\phi(w) = \text{true}$ **then**
        $S' := S' \cup \{w\}$
      end
    end
    $\mathbf{S}' := \mathbf{S}' \cup \{S'\}$;
  end
end

Table 2a. Strategy I_0: Pruning by keeping only heads of sentences

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	38.54	57.46	26.88	17.43	31.06	88.49	51.90	44.54
Precision	49.13	66.18	49.61	65.73	39.07	42.18	60.31	53.17
F-measure	43.19	61.51	34.87	27.55	34.61	57.13	55.79	44.95

MinSupp=0.004, MinConf=36.6, Var.=2.14, Disp.=47.84, AvgTransSize=1.00

Table 2b. Strategy I_1: Pruning by keeping only nsubj $\to$ head dependencies

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	70.55	60.44	66.97	58.27	63.22	78.76	69.92	66.88
Precision	73.05	80.84	72.87	86.98	71.71	85.29	76.29	78.15
F-measure	71.78	69.17	69.80	69.79	67.19	81.89	72.97	71.80

MinSupp=0.007, MinConf=60.4, Var.=1.43, Disp.=41.71, AvgTransSize=1.04

a given morphosyntactic constraint ϕ. The following strategies correspond to various definitions of ϕ.

Strategy $\mathbf{I_0}$ Our first strategy will be to keep only the head of each sentence (which, incidentally, is a verb in 85.37% of sentences of our corpus). This corresponds to the constraint $\phi(w) \equiv (w = \text{root}(S))$. Results are given on Table 2a.

Although the recall of GSPO is quite high (a possible interpretation could be that sports use very specific verbs), F-measure is quite low when we compare it to the one of the tfidf-based method of the same average itemset length, namely 65.69%.

Strategy $\mathbf{I_1}$ The second strategy consists in keeping words connected to the head by a (single) dependency of type nsubj (= nominal subject). This occurs in

Table 2c. Strategy I'_1: Pruning by keeping only ccomp → head dependencies

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	57.33	33.82	25.31	16.96	21.74	47.62	59.60	37.48
Precision	37.83	47.55	38.50	42.06	34.62	57.05	54.17	44.54
F-measure	45.59	39.53	30.54	24.17	26.71	51.91	56.75	39.31

MinSupp=0.008, MinConf=34.4, Var.=1.97, Disp.=19.59, AvgTransSize=1.15

Table 2d. Strategy I_2: Pruning by keeping only nouns at distance 1 from head

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	80.75	75.92	73.24	68.59	70.59	95.55	77.96	77.51
Precision	73.21	83.51	75.35	89.86	73.67	80.52	77.36	79.07
F-measure	76.80	79.53	74.28	77.80	72.09	87.39	77.66	77.94

MinSupp=0.016, MinConf=51.6, Var.=2.43, Disp.=244.82, AvgTransSize=2.70

Table 2e. Strategy I'_2: Pruning by keeping only verbs at distance 1 from head

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	54.32	62.50	44.37	20.41	27.58	91.39	67.68	52.61
Precision	49.58	65.98	48.44	78.39	46.69	43.57	63.84	56.64
F-measure	51.84	64.19	46.32	32.39	34.68	59.01	65.70	50.59

MinSupp=0.019, MinConf=30, Var.=4.00, Disp.=175.41, AvgTransSize=2.01

79.84% of sentences of our corpus. The constraint is then $\phi(w) \equiv (\exists(w, \text{root}(S), \text{nsubj}) \in \text{Dep}(S))$. Results are given on Table 2b.

Note that the slightly higher than 1 transaction size is probably due to the rare cases where there are more than one nsubj dependencies pointing to the head. The scores rise dramatically when compared to those of the strategy based only on the head of the sentence. The average F-measure (71.80%) is significantly higher than the tfidf-based performance for the same average transaction size (65.69%). *This shows that using a dependency property to select a word is a better choice than the one provided by the frequentist tfidf-based method.* Note that the case of nsubj is unique: if we take ccomp (= clausal complement) instead of nsubj, the performance falls even below the level of strategy I_0 (Table 2c).

Strategy I_2 The third strategy considers all nouns (POS tags starting with `N`) at distance 1 from the head in the dependency graph. Such dependencies occur in 59.24% of the sentences of our corpus. This corresponds to $\phi(x) \equiv ((\exists(x, \text{root}(S), d) \in \text{Dep}(S)) \wedge (\text{POS}(x) = \texttt{N*}))$. Results are given on Table 2d.

The result seems better than the one of strategy I_1 (Table 2b). However, if we take transaction size into account, it is in fact merely equivalent to—and hence not better than, as it was the case for I_1—the tfidf-based method with the same transaction size. Once again we see a very high recall rate for the sports category.

Algorithm 6: Corpus hyperonymization

Data: A dependency-pruned corpus $\mathbf{S}'$, an hyperonymic function MSCH: $\mathcal{W} \to \mathcal{W}^{\mathbb{N}}$, the hyperonymic order N
Result: The hyperonymically extended corpus $\mathbf{S}''$
`Hyperonymize(`$\mathbf{S}'$, MSCH, N`):`

```
S'' := ∅;
for S' ∈ S' do
    S'' := ∅;
    for w ∈ S' do
        if proj_N(MSCH(w)) ≠ ∅ then
            S'' := S'' ∪ {proj_N(MSCH(w))}
        else
            S'' := S'' ∪ {w}
        end
    end
    S'' := S'' ∪ {S''}
end
end
```

One could be tempted to check the performance of taking verbs (instead of nouns) at distance 1 from the head. Indeed, verbs at that position are more frequent than nouns: they occur in 62.94% of the sentences of our corpus. Nevertheless, the results are not as good (Table 2e). This shows that despite their high frequency, verbs contain less pertinent information than nouns at the same distance from the head.

3.4 Methods based on dependencies and hyperonyms

In this section we add semantic information by the means of hyperonyms, using the hyperonymization function h (§ 1.3). The preprocessing is done by Alg. 6: $h_i(w)$ is an N-th order hyperonym of w, if it exists in WordNet. In case there is no N-th order hyperonym, the word remains unchanged. We call N the *hyperonymic factor* of our itemset transformation.

Strategy $\mathbf{II_1}$ This strategy considers hyperonymic factor $N = 1$. We thus first apply strategy I_1 and then hyperonymization h_1. Results are presented on Table 3a.

The performance is globally inferior to the one of Strategy I_1 (in which, F-measure attained 71.80%). It is interesting to note that the recall of class GJOB has decreased significantly (48.96% vs. 58.27%): in other words, using hyperonyms when dealing with labor issues results into failure to recognize 9.31% of the documents as belonging to the domain; one could say that terms used in GJOB lose their "labor specificity" already at first-order hyperonymization. On the other hand, the (already high in I_1) recall of GSPO has increased even more, compared to I_1 (from 78.76% to 82.20%): it seems that sports terminology remains in the domain even after hyperonymization, and replacing specific terms

Table 3a. Strategy II_1: I_1 followed by first-order hyperonymization

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	72.39	56.04	71.32	48.96	59.02	82.20	70.42	65.76
Precision	66.21	75.42	64.33	82.21	67.71	73.20	70.73	71.40
F-measure	69.16	64.30	67.64	61.37	63.07	77.44	70.57	67.65

MinSupp=0.010, MinConf=44.8, Var.=1.92, Disp.=78.47, AvgTransSize=1.04

Table 3b. Strategy II_2: I_1 followed by second-order hyperonymization

	E12	GCRIM	GDIP	GJOB	GPOL	GSPO	GVIO	AVG
Recall	69.02	52.78	71.97	47.77	54.24	80.80	65.67	63.18
Precision	64.94	74.57	61.23	80.64	65.81	69.96	72.33	69.93
F-measure	66.92	61.81	66.16	60.00	59.47	74.99	68.84	65.46

MinSupp=0.008, MinConf=44.8, Var.=1.85, Disp.=70.51, AvgTransSize=1.04

by more general ones has increased their frequency as items, and hence improved recall. We have the same phenomenon with the recall of GDIP (which increased from 66.97% to 71.32%), and also slightly with the recalls of E12 and GVIO.

Strategy II_2 This strategy is similar to strategy II_1 but uses hyperonymic factor $N = 2$. Results are presented on Table 3b.

The performance is globally inferior to the one of II_1 (where we used first-order hyperonyms), with two minor exceptions: the recall of GDIP that increased by 0.65% and the precision of GVIO that increased by 1.6%. What is noteworthy however, is the fact that the recalls of GDIP and GSPO are still higher than the ones of strategy I_1 (no hyperonyms).

To better understand the behavior of the system when climbing the hyperonymic chain by replacing words by hyperonyms of increasingly higher order (and returning to the original word when there are no hyperonyms left) we calculated the performance for N-th order hyperonyms for $1 \leq N \leq 12$. Note that when $N > 12$ the amount of remaining hyperonyms is negligible and the strategy is similar to strategy I_1 (no hyperonyms). On Fig. 2, the reader can see the evolution of recall (black), precision (red) and F-measure (blue) for the average of all class, and then specifically for GSPO and for GDIP. Dashed lines represent the recall, precision and F-measure of strategy I_1.

In the average case, the effect of hyperonymization of orders 1–4 is to decrease performance. After $N = 5$, the global number of hyperonyms available in WordNet rapidly decreases so that the situation gradually returns to the one of I_1 (no hyperonyms) and we see curves asymptotically converging to I_1 lines from underneath.

Not so for GSPO, the GSPO recall curve of which is above the I_1 value for most N ($N = 1$, 2, 6–8 and 10–12).

The phenomenon is even better illustrated in the case of GDIP: as the reader can see on the figure, the *complete GDIP recall curve is located above the I_1 one*. It seems that in these two cases (GDIP and, to a lesser extent, GSPO),

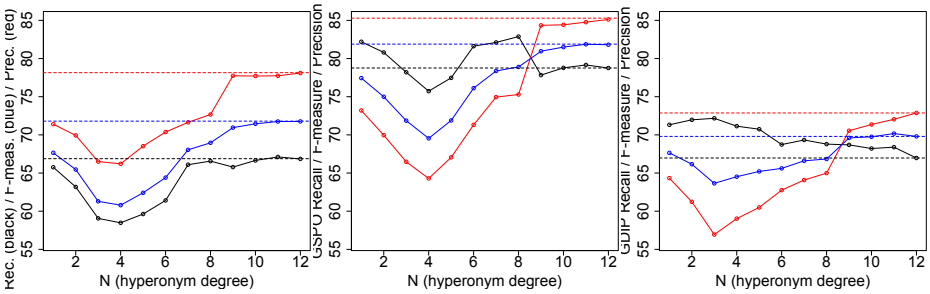
Rec. (black) / F-meas. (blue) / Prec. (red)
GSPO Recall / F-measure / Precision
GDIP Recall / F-measure / Precision
N (hyperonym degree)
N (hyperonym degree)
N (hyperonym degree)

Fig. 2. F-1 measure for hyperonymization of orders $1 \leq N \leq 12$: the average case, class GSPO, class GDIP

hyperonyms of all orders have a positive impact on the classifier. Unfortunately this impact only concerns recall and is compensated by bad precision, so that F-measure is still inferior to the I_1 case.

4 Conclusion and future work

In this paper we have investigated the use of association rules for text classification by applying two new techniques: (a) we reduce the number of word features through the use of morphosyntactic criteria in the framework of dependency syntax; for that we keep words dependent from the head by specific dependencies and/or having specific POS tags (b) we replace words by their hyperonyms of different orders, which we have calculated out of WordNet using frequencies and, in some cases, disambiguation. We have obtained positive results for case (a), in particular when we compare dependency-based single-item rules with tfidf-based ones. In case (b) the results we share in this paper are less efficient but still interesting, especially we found classes for which hyperonymization significantly improves recall.

This work opens several perspectives, among which:

— examine why these particular classes are favorable to hyperonymization, whether this is related to the structure of WordNet or to linguistic properties of the domain;

— explore *partial hyperonymization* i.e., is it possible to hyperonymize only specific items according to the needs of the classifier?[1] How do we choose, on the word level, if we should rather keep the original word (to increase precision) or switch to some hyperonym (to increase recall)?

— we have used only recall and precision as quality measures of our rules, and our evaluation is strongly dependent on these measures since the selection of the 1,000 rules we keep is entirely based upon them. There are other quality

[1] Indeed, by hyperonymizing all words one wins on one side and loses on the other: for example “dalmatian” and “poodle” will both be replaced by “dog”, but “dog” occurrences will be replaced by “canid”. It would be more preferable to keep the word “dog” in the second case, so that we have a real increase in frequency.

measures available, how do they apply and how can they be compared and combined? How robust are the results?

— and finally: how can we optimize the distinctive feature of association rules, namely the fact of being intelligible by the user? How can the user's experience (and linguistic knowledge) be incorporated in the enhancement of rules to obtain the best possible result from his/her point of view?

References

1. Reuters corpus, volume 1, english language, 1996-08-20 to 1997-08-19, `http://about.reuters.com/researchandstandards/corpus/statistics/index.asp`
2. British National Corpus (1994), `http://www.natcorp.ox.ac.uk`
3. Aggarwal, C.C., Zhai, C.: A Survey of Text Classification Algorithms, chap. 6, pp. 163–222. Mining Text Data, Springer
4. Ahonen, H., Heinonen, O., Klemettinen, M., Verkamo, A.I.: Applying data mining techniques in text analysis. TR C-1997-23, Department of Computer Science, University of Helsinki (1997)
5. Borgelt, C.: Efficient implementations of apriori and eclat. In: Workshop of Frequent Item Set Mining Implementations (FIMI 2003), Melbourne, FL, USA (2003)
6. Ferrer i Cancho, R., Solé, R.V., Köhler, R.: Patterns in syntactic dependency networks. Physical Review E 69, 1–8 (2004)
7. Cherfi, H., Napoli, A., Toussaint, Y.: A Conformity Measure Using Background Knowledge for Association Rules: Application to Text Mining, pp. 100–115. IGI Global (2009)
8. Chomsky, N.: Syntactic structures. Mouton (1957)
9. Cohen, W.W.: Learning rules that classify e-mail. In: AAAI Spring Symposium on ML and IR (1996)
10. Curran, J.R., Moens, M.: Scaling context space. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics. pp. 231–238 (2002)
11. Do, T.N.Q.: A graph model for text analysis and text mining. Master Thesis, Université de Lorraine, (2012)
12. Jaillet, S., Laurent, A., Teisseire, M.: Sequential patterns for text categorization. Intell. Data Anal. 10(3), 199–214 (2006)
13. Kovacs, L., Baksa-Varga, E.: Dependency-based mapping between symbolic language and Extended Conceptual Graph. In: 6th International Symposium on Intelligent Systems and Informatics. pp. 1–6 (2008)
14. Lang, K.: Newsweeder: Learning to filter netnews. In: International Conference on Machine Learning. pp. 331–339. Morgan Kaufmann (1995)
15. Liu, B., Hsu, W., Ma, Y.: Integrating classification and association rule mining. In: Proc. of the Int. Conf. on Knowledge Discovery and Data Mining (New York). pp. 80–86 (1998)
16. Lowe, W.: Towards a theory of semantic space. In: Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society. pp. 576–581 (2001)
17. de Marneffe, M.C., MacCartney, B., Manning, C.D.: Generating typed dependency parses from phrase structure parses. In: LREC 2006. pp. 449–454 (2006)
18. Mehran Sahami, Susan Dumais, D.H., Horvitz, E.: A bayesian approach to filtering junk email. In: AAAI Workshop on Learning for Text Categorization. AAAI Technical Report WS-98-05 (1998)

19. Mel′čuk, I.A.: Dependency syntax : theory and practice. Albany: State University Press of New York (1987)
20. Miller, G.A.: WordNet: A lexical database for English. Communications of the ACM 38(11), 39–41 (1995)
21. Nivre, J.: Dependency Grammar and Dependency Parsing. MSI report 05133. School of Mathematics and Systems Engineering, Växjö University, (2005)
22. Ordoñez-Salinas, S., Gelbukh, A.: Information Retrieval with a Simplified Conceptual Graph-Like Representation 6437(Chapter 9), 92–104 (2010)
23. Padó, S., Lapata, M.: Dependency-based construction of semantic space models. Computational Linguistics 33(2), 161–199 (2007)
24. Pedersen, T.: Information content measures of semantic similarity perform better without sense-tagged text. In: Proceedings of the 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL HLT 2010). pp. 329–332 (2010)
25. Roche, M., Azé, J., Matte-Tailliez, O., Kodratoff, Y.: Mining texts by association rules discovery in a technical corpus. In: Intelligent Information Processing and Web Mining. pp. 89–98. Advances in Soft Computing, Springer Verlag (2004)
26. Schmid, H.: Probabilistic part-of-speech tagging using decision trees. In: Proceedings of International Conference on New Methods in Language Processing, Manchester, UK (1994)
27. Sebastiani, F.: Machine learning in automated text categorization. ACM Comput. Surv. 34(1), 1–47 (2002)
28. Tesnière, L.: Éléments de syntaxe structurale. Klincksieck (1959)
29. Zaïane, O.R., Antonie, M.L.: Classifying text documents by associating terms with text categories. In: Australasian Database Conference. CRPIT, vol. 5. Australian Computer Society (2002)